

ff



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/782,532	02/13/2001	Gregory Hagan Moulton	UND011	7507

7590

10/31/2005

Stuart T. Langley, Esq.  
Hogan & Hartson, LLP  
Suite 1500  
1200 17th Street  
Denver, CO 80202

EXAMINER
----------

PHILLIPS, HASSAN A

ART UNIT	PAPER NUMBER
----------	--------------

2151

DATE MAILED: 10/31/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

**Office Action Summary**

Application No.

09/782,532

Applicant(s)

MOULTON ET AL.

Examiner

Hassan Phillips

Art Unit

2151

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --  
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

**Status**

- 1) ☒ Responsive to communication(s) filed on 22 August 2005.  
2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.  
3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

**Disposition of Claims**

- 4) ☒ Claim(s) 1,3-10,12-21,23,24,26-28,33-37 and 39-54 is/are pending in the application.  
4a) Of the above claim(s) \_\_\_\_\_ is/are withdrawn from consideration.  
5) ☐ Claim(s) \_\_\_\_\_ is/are allowed.  
6) ☒ Claim(s) 1,3-10,12-21,23,24,26-28,33-37 and 39-54 is/are rejected.  
7) ☐ Claim(s) \_\_\_\_\_ is/are objected to.  
8) ☐ Claim(s) \_\_\_\_\_ are subject to restriction and/or election requirement.

**Application Papers**

- 9) ☐ The specification is objected to by the Examiner.  
10) ☐ The drawing(s) filed on \_\_\_\_\_ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.  
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).  
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).  
11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

**Priority under 35 U.S.C. § 119**

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).  
a) ☐ All b) ☐ Some \* c) ☐ None of:  
1. ☐ Certified copies of the priority documents have been received.  
2. ☐ Certified copies of the priority documents have been received in Application No. \_\_\_\_\_.  
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

\* See the attached detailed Office action for a list of the certified copies not received.

**Attachment(s)**

- 1) ☒ Notice of References Cited (PTO-892)  
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)  
3) ☒ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)  
Paper No(s)/Mail Date 8/22/05.  
4) ☐ Interview Summary (PTO-413)  
Paper No(s)/Mail Date. \_\_\_\_\_.  
5) ☐ Notice of Informal Patent Application (PTO-152)  
6) ☐ Other: \_\_\_\_\_.

### **DETAILED ACTION**

1. This action is in response to communications filed on August 22, 2005.

#### ***Response to Arguments***

2. Applicant's arguments filed August 22, 2005 have been fully considered but they are not persuasive. Applicant argued that:

- a) Carter does not show "network connected devices that...are located at distinct network nodes from the plurality of network-accessible storage devices";
- b) The teachings of Carter contrast with Applicants claim 21 since claim 21 calls for "storage management process instances that are distributed across the network-accessible devices such that failure or unavailability of any given instance of a storage management process instance will not impact the availability of stored data";
- c) Carter fails to show or suggest any recognition of migrating data in response to a determination that a fault condition is likely;
- d) The statements cited at page 3, lines 10-15 of the application are not fairly attributed to the work of "another" as required by MPEP 2129;
- e) Carter does not state that replication is done at different nodes;
- f) Replicating data is not the same as RAID-type distribution as called for in claim 1;

- g) Carter does not show or suggest the N-dimensional parity scheme of claims 12-15 or greater than two dimensions of parity as called for in claim 18;
- h) Carter alone or in combination with Litwin does not show identifying two or more storage devices at different network locations,... determining parity data for the data to be stored,... and storing the data and/or parity data using a raid type distribution across the two or more storage devices;
- i) Gershman does not supply the basic deficiency of Carter in that Gershman does not teach or suggest monitoring a plurality of network accessible storage devices using the heartbeat message;
- j) McClain does not show or suggest "storage management processes are configured to migrate data amongst the storage devices using the storage messages preemptively when a fault condition in at least one of the storage devices is determined to be likely.";
- k) Thompson does not show or suggest a method of data storage management that calls for "storage management process instances...distributed across the network-accessible devices such that failure or unavailability of any given instance of a storage management process instance will not impact the availability of stored data."

Examiner respectfully disagrees with Applicant's assertions.

3. Regarding item a), Examiner acknowledges the teachings of Carter disclose a shared memory system where each node implements a portion of the shared memory. Examiner also recognizes that the system disclosed by Carter is not limited to these teachings. In fact, Carter teaches providing "a global addressable storage system that allows remote computers and computers on different, interconnected networks to communicate and share data in a transparent and dynamic manner", (col. 3, lines 1-5). Carter further teaches "a computer system comprising a first computer network coupled to a second, remote computer network. The first computer network includes a first plurality of computers, a first persistent storage device, and a first globally addressable data storage system that maintains and allows access to data on the first network and that provides addressable access to data stored in the first persistent data storage device. The second computer network includes a second plurality of computers, a second persistent data storage device, and a second global addressable data storage system that maintains and allows access to data on the second network and that provides addressable access to data stored in the second persistent data storage device. The first and second globally addressable data storage systems interoperate to allow the first computers to access data on the second network including data stored in the second persistent data storage device and to allow the second computers to access data on the first network including data stored in the first persistent data storage device.", (col. 4, lines 15-37). Clearly these teaching of Carter suggest "network connected devices that...are located at distinct network nodes from the plurality of network-accessible storage devices" as claimed by Applicant.

4. Regarding item b), Examiner acknowledges the teachings of Carter disclose failure of any shared memory subsystem will render the portion of the shared memory allocated to that node and the data stored therein unavailable. Carter further teaches however, storage management process instances distributed across the nodes such that failure or unavailability of any given node will not impact the availability of stored data since the data is replicated to other nodes and may be accessed therefrom, (col. 7, lines 43-60, col. 8, lines 31-50, col. 23, lines 12-27). Thus, Applicant's claimed invention fails to teach away from the teachings of Carter since Applicant's claimed invention fails to disclose whether or not the stored data could be accessed from more than one node upon failure or unavailability of any given instance of a storage management process.

5. Regarding item c), as indicated in previous actions, it is inherent in the teachings of Carter that data is migrated when a fault condition is determined to be likely, (col. 23, lines 12-27). Applicant's claimed invention recites, "...wherein the storage management processes are configured to migrate data amongst the storage devices using the storage messages preemptively when a fault condition in at least one of the storage devices is determined to be likely." The claimed invention fails to expressly indicate whom, or what is doing the determining. Thus, giving broadest reasonable interpretation to the claimed invention, Carter reads over the claimed invention because, for example, a user of Carter's system would include a coherent

Art Unit: 2151

replication controller to provide for fault tolerant operation of the shared memory subsystem because the user has anticipated fault conditions in at least one of the storage devices was determined to be likely, (col. 23, lines 12-27).

6. Regarding item d), as indicated in previous actions, the teachings of Carter disclose substantial features of Applicant's claimed invention. Features of Applicant's claimed invention that were not expressly disclosed by Carter (i.e. RAID-type distributions) were well known in the art and suggested in the teachings of Carter, (col. 8, lines 42-50). Examiner respectfully submits that the lines cited in Applicant's application are lines disclosing "conventional" teachings that were not expressly disclosed by Carter. Applicant's recently provided reference, Dziadosz, U.S. 5,832,222 further provides evidence that such features were well known in the art at the time of the present invention, (Dziadosz, col. 3, lines 50-66).

7. Regarding item e) it is unclear to the Examiner why Applicant believes replication is not done at different nodes in the teachings of Carter. In previous actions Examiner has indicated Carter teaches maintaining "data coherence among network nodes; automatically replicates data for redundancy and fault tolerance; automatically and dynamically migrates data to account for varying network usage and traffic patterns", (col. 8, lines 42-50). Clearly this suggests replication is done at different nodes.

8. Regarding item f) amended claim 1 recites, "...wherein the processes for storing data comprise processes that implement a RAID-type distribution across the plurality of network-accessible devices such that write operations comprise writes to multiple network nodes and read operations comprise reads from multiple network nodes". As indicated in previous actions Examiner has interpreted a "RAID-type distribution" as a storage method in which data is distributed across a group of computers on the network that function as a single unit. The amended claims fail to expressly indicate that a "RAID-type distribution" involves executing every data write operation to a primary node and all mirror nodes as set out at page 28, lines 9-11 in the application. Applicant is reminded that although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

9. Regarding item g) Examiner submits the combined teachings of Carter and AAPA provide a means for the N-dimensional parity scheme of claims 12-15 or greater than two dimensions of parity as called for in claim 18 for reasons indicated above and in previous actions.

10. Regarding item h) Examiner submits the combined teachings of Carter and Litwin provide a means identifying two or more storage devices at different network locations,...determining parity data for the data to be stored,...and storing the data



Art Unit: 2151

and/or parity data using a raid type distribution across the two or more storage devices for reasons indicated above and in previous actions.

11. Regarding item i) Examiner submits the combined teachings of Carter and Gershman provide a means for monitoring a plurality of network accessible storage devices using the heartbeat message for reasons indicated in previous actions. In response to applicant's argument that there is no suggestion to combine the references, the examiner recognizes that obviousness can only be established by combining or modifying the teachings of the prior art to produce the claimed invention where there is some teaching, suggestion, or motivation to do so found either in the references themselves or in the knowledge generally available to one of ordinary skill in the art. See *In re Fine*, 837 F.2d 1071, 5 USPQ2d 1596 (Fed. Cir. 1988) and *In re Jones*, 958 F.2d 347, 21 USPQ2d 1941 (Fed. Cir. 1992). In this case, it would have been obvious to one of ordinary skill in the art to modify the teachings of Carter to show monitoring the data storage for faults by means of the plurality of storage management processes, wherein the monitoring comprises at least a portion of the plurality of network accessible storage devices transmitting heartbeat messages since doing so would have provided an efficient means for the storage management processes to take appropriate action when a fault is detected through the monitoring, Gershman, col. 47, lines 65-67, col. 48, lines 1-4.

12. Regarding item j) Examiner submits Carter discloses storage management processes are configured to migrate data amongst the storage devices using the storage messages preemptively when a fault condition in at least one of the storage devices is determined to be likely for reasons indicated above.

13. Regarding item k) Examiner submits Carter discloses storage management process instances...distributed across the network-accessible devices such that failure or unavailability of any given instance of a storage management process instance will not impact the availability of stored data for reasons indicated above.

14. Accordingly the references supplied by the examiner in the previous office action covers the claimed limitations. The rejections are thus sustained. Applicant is requested to review the prior art of record for further consideration.

***Claim Rejections - 35 USC § 102***

15. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) The invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Art Unit: 2151

16. Claims 21, 23, 24, 27, 28, 44-46, 48-50, are rejected under 35 U.S.C. 102(e) as being anticipated by Carter et al. (hereinafter Carter), U.S. patent 5,987,506.

17. In considering claim 21, Carter discloses a method of data storage management comprising the acts of: providing a plurality of network-accessible storage devices capable of storing data, some of which are located at distinct network nodes, (col. 3, lines 1-5, col. 4, lines 15-37, col. 6, lines 7-12); providing a plurality of network-connected devices that access the network-accessible storage devices via the network, (col. 3, lines 1-5, col. 4, lines 15-37, col. 6, lines 7-12); wherein the plurality of network-connected devices are located at distinct network nodes from the plurality of network-accessible storage devices, (col. 3, lines 1-5, col. 4, lines 15-37, col. 6, lines 7-12); implementing a plurality of storage management process instances wherein the storage management process instances are distributed across the network-accessible devices such that failure or unavailability of any given instance of a storage management process instance will not impact the availability of stored data, (col. 7, lines 43-60, col. 8, lines 31-50, col. 23, lines 12-27); communicating storage messages between the storage management process instances, (col. 7, lines 43-60); storing data to the network-accessible device under control of at least one instance of the storage management processes, (col. 7, lines 43-60); and implementing a peer-to-peer network between the plurality of network-accessible storage devices, (col. 13, lines 19-58); communicating state information for the plurality of network-accessible storage devices between the plurality of network-accessible storage devices, (col. 13, lines 19-58); and

Art Unit: 2151

performing read and write operations using the plurality of storage devices such that write operations comprise writes to multiple network nodes and read operations comprise reads from multiple network nodes, (col. 7, lines 43-60, col. 13, lines 19-58).

18. In considering claim 23, the method disclosed by Carter teaches serving data from the plurality of network accessible storage device. See col. 7, lines 43-49.

19. In considering claim 24, the method disclosed by Carter further teaches the plurality of storage device comprising a RAID storage system. See col. 16, lines 43-46.

20. In considering claim 27, the method disclosed by Carter teaches the processes for storing data comprising processes that provide a means for storing parity and/or mirror data across more than one of the plurality of network accessible storage devices. See col. 4, lines 10-38, and col. 8, lines 39-50.

21. In considering claim 28, the method disclosed by Carter further teaches the storage management processes further comprising processes for recovery of data when one or more of the plurality of network accessible storage devices is unavailable. See col. 36, lines 10-22.

22. In considering claim 44, Carter discloses a method of data storage management comprising the acts of: providing a plurality of network-accessible storage

Art Unit: 2151

devices capable of storing data, (col. 6, lines 37-43); implementing a plurality of storage management process instances to communicate with each other, (col. 7, lines 43-60); the processes being capable of storage allocation and de-allocation across the plurality of network-accessible storage devices, (col. 21, lines 32-44); wherein the storage management processes are inherently configured to migrate data amongst the storage devices using the storage messages preemptively when a fault condition in at least one of the storage devices is determined to be likely. (col. 23, lines 12-27).

23. In considering claim 45, the disclosed method of Carter teaches the processes being configured to use the storage messages to reconstruct data stored in a failed one of the storage devices. See col. 36, lines 10-63.

24. In considering claim 46, the disclosed method of Carter teaches the storage management processes being configured to migrate data amongst the storage devices using the storage messages in response to a detected fault condition in at least one of the storage devices. See col. 23, lines 12-27.

25. In considering claim 48, the disclosed method of Carter teaches the plurality of storage devices comprising an arbitrarily large number of storage devices. See col. 9, lines 10-17.

Art Unit: 2151

26. In considering claim 49, the disclosed method of Carter provides a means for associating parity information with a data set, storing the parity information in at least some of the storage devices, and serving data requests corresponding to the data set by accessing the parity information associated with the data set. See col. 23, lines 12-27.

27. In considering claim 50, the disclosed method of Carter teaches storing a data set in a plurality of the data storage devices using the storage management processes. Further, it is inherent in the method disclosed by Carter that data requests corresponding to the data set are accessed from the plurality of data storage devices in parallel. See col. 7, lines 8-38.

***Claim Rejections - 35 USC § 103***

28. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

29. Claims 1, 3-10, 12-20, 26, 52, 53, are rejected under 35 U.S.C. 103(a) as being unpatentable over Carter, in view of Applicants Admitted Prior Art (AAPA).

30. In considering claim 1, Carter discloses a data storage management system comprising: a plurality of network-accessible storage devices capable of storing data, some of which are located at distinct network nodes, (col. 6, lines 7-12); a plurality of network-connected devices that access the network-accessible storage devices via the network, (col. 3, lines 1-5, col. 4, lines 15-37, col. 6, lines 7-12); wherein the plurality of network-connected devices are located at distinct network nodes from the plurality of network-accessible storage devices, (col. 3, lines 1-5, col. 4, lines 15-37, col. 6, lines 7-12); a plurality of network-accessible devices configured to implement storage management processes wherein the storage management processes are distributed across the network-accessible devices such that failure or unavailability of any given instance of a storage management process will not impact the availability of stored data, (col. 7, lines 43-60, col. 8, lines 31-50, col. 23, lines 12-27); a communication system enabling the storage management processes to communicate with each other, (col. 7, lines 43-60); wherein the storage management processes comprise processes for storing data to the at least one network-accessible device, (col. 7, lines 43-60); and performing read and write operations using the plurality of storage devices such that write operations comprise writes to multiple network nodes and read operations comprise reads from multiple network nodes, (col. 7, lines 43-60, col. 13, lines 19-58).

Although the disclosed system and method taught by Carter shows substantial features of the claimed invention, it fails to expressly disclose: the storage management processes implementing a RAID-type distribution across the plurality of network-accessible devices.

Nevertheless, the teachings of Carter provide, and suggest, a means for implementing a RAID-type distribution across the plurality of network devices, where Carter discloses: the storage management processes maintaining data coherence among the network nodes, and automatically replicating data for redundancy and fault tolerance, (col. 8, lines 42-50). Furthermore, as admitted by the Applicant in the Applicant's disclosure of the claimed invention, RAID-type systems were well known in the art at the time of the present invention, (page 3, lines 7-10).

Thus, if not implicit in the teachings of Carter, it would have been obvious to a person of ordinary skill in the art at the time of the present invention to modify the teachings of Carter to show the storage management processes implementing a RAID-type distribution across the plurality of network accessible devices. Doing so would have improved performance of the storage management system taught by Carter, and would have also increased fault tolerance, (AAPA, page 3, lines 10-15).

31. In considering claim 3, the system disclosed by Carter teaches the storage management processes comprising processes for serving data from the plurality of network accessible storage devices. See col. 7, lines 43-49.

32. In considering claim 4, the system disclosed by Carter further teaches the plurality of storage devices comprising a RAID storage system. See col. 16, lines 43-46.



33. In considering claim 5, the system disclosed by Carter teaches the plurality of network accessible storage devices comprising a computer with the direct attached storage (DAS) being any persistent memory. This provides a means for the DAS to be selected from the group consisting of magnetic hard disk, magneto-optical, optical disk, digital optical tape, holographic storage, quantum storage, and atomic force probe storage. See col. 16, lines 43-46.

34. In considering claim 6, the system disclosed by Carter further teaches the plurality of network-accessible storage devices comprising a peer-to-peer network of storage devices, each network-accessible storage device having means for communicating state information with other network-accessible storage devices, at least one network-accessible storage device comprising means for receiving storage requests from external entities, and at least one network-accessible storage device comprising means for causing read and write operations to be performed on others of the network-accessible storage devices. See col. 7, lines 43-60 and col. 13, lines 19-58.

35. In considering claim 7, the system disclosed by Carter provides a means for the communication system to comprise a TCP/IP over Ethernet network. See col. 3, lines 1-5.

Art Unit: 2151-

36. In considering claim 8, the system disclosed by Carter provides a means for the communication system to comprise a Gigabit Ethernet network. See col. 3, lines 1-5.

37. In considering claim 9, the system disclosed by Carter provides a means for the communication system to comprise a Fiber Channel fabric. See col. 3, lines 1-5.

38. In considering claim 10, the system disclosed by Carter provides a means for the communication system to comprise a wireless network. See col. 3, lines 1-5.

39. In considering claim 12, although the disclosed system and method taught by Carter shows substantial features of the claimed invention, it fails to expressly disclose: the storage management processes implementing an n-dimensional parity scheme for data elements across the plurality of network-accessible storage devices.

Nevertheless, the teachings of Carter provide, and suggest, a means for implementing an n-dimensional parity scheme for data elements across the plurality of network-accessible storage devices, where Carter discloses: the storage management processes maintaining data coherence among the network nodes, and automatically replicating data among the network devices, (col. 4, lines 10-38, col. 8, lines 42-50). Furthermore, as admitted by the Applicant in the Applicant's disclosure of the claimed invention, parity schemes were well known in the art at the time of the present invention, (page 5, lines 28-31).

Thus, if not implicit in the teachings of Carter, it would have been obvious to a person of ordinary skill in the art at the time of the present invention to modify the teachings of Carter to show the storage management processes implementing an n-dimensional parity scheme across the plurality of network accessible devices. Doing so would have improved performance of the storage management system and method taught by Carter, and would have also increased fault tolerance, (AAPA, page 3, lines 10-15).

40. In considering claim 13, the teachings of Carter provides a means for expanding or contracting the size of "n" in the n-dimensional parity scheme for the data to be stored by the plurality of network accessible devices to whatever extent is desired. See Carter, col. 4, lines 10-38 and col. 8, lines 39-50.

41. In considering claim 14, the system disclosed by Carter further teaches the storage management processes further comprising processes for recovery of data when one or more of the network accessible storage devices is unavailable. See col. 36, lines 10-22.

42. In considering claim 15, the system disclosed by Carter teaches the storage management processes comprising processes that provide a means for accessing stored data when one or more of the network accessible storage devices are not

Art Unit: 2151

desirable data sources for reasons including efficiency, performance, network congestion, and security. See col. 6, lines 12-14.

43. In considering claim 16, the system disclosed by Carter further provides a means for the plurality of network-accessible devices configured to implement storage management processes to further comprise commercial off-the-shelf computer systems implementing a common operating system. See col. 18, lines 63-67, col. 19, lines 1-4.

44. In considering claim 17, the system disclosed by Carter further provides a means for the plurality of network-accessible devices configured to implement storage management processes to further comprise commercial off-the-shelf computer systems implementing a heterogeneous set of operating systems. See col. 18, lines 63-67, col. 19, lines 1-4.

45. In considering claim 18, the teachings of Carter provides a means for the storage management processes to comprise processes for implementing greater than two dimensions of parity. See Carter, col. 4, lines 10-38, and col. 8, lines 39-50.

46. In considering claim 19, the system disclosed by Carter teaches the processes for storing data comprising processes that provide a means for storing parity and/or mirror data across more than one of the plurality of network accessible storage devices. See col. 4, lines 10-38, and col. 8, lines 39-50.

47. In considering claim 20, the teachings of Carter suggest the storage management processes comprise processes for adding and removing additional storage capacity to individual network-accessible storage devices and the system as a whole. See col. 9, lines 10-17.

48. In considering claim 26, although the disclosed method taught by Carter shows substantial features of the claimed invention, it fails to expressly disclose: the step of storing data comprising storing data using a RAID-type distribution across the plurality of network-accessible devices.

Nevertheless, the teachings of Carter provide, and suggest, a means for storing data using a RAID-type distribution across the plurality of network devices, where Carter discloses: the storage management processes maintaining data coherence among the network nodes, and automatically replicating data for redundancy and fault tolerance, (col. 8, lines 42-50). Furthermore, as admitted by the Applicant in the Applicant's disclosure of the claimed invention, RAID-type systems were well known in the art at the time of the present invention, (page 3, lines 7-10).

Thus, if not implicit in the teachings of Carter, it would have been obvious to a person of ordinary skill in the art at the time of the present invention to modify the teachings of Carter to show the step of storing data comprising storing data using a RAID-type distribution across the plurality of network accessible devices. Doing so

Art Unit: 2151

would have improved performance of the storage management system taught by Carter, and would have also increased fault tolerance, (AAPA, page 3, lines 10-15).

49. In considering claim 52, the teachings of Carter provide a means for RAID-type distribution comprising managing redundancy operations across the plurality of network-accessible devices. See col. 8, lines 31-50.

50. In considering claim 53, the teachings of Carter provide a means for RAID-type distribution comprising one or more functionalities selected from the group consisting of data striping, data mirroring, parity data distribution, parity data mirroring, and data entry as N-separated secrets. See col. 23, lines 12-27.

51. Claims 33-37, 39-41, are rejected under 35 U.S.C. 103(a) as being unpatentable over Carter in view of Litwin et al. (hereinafter Litwin), U.S. patent 6,122,754 (supplied by applicant) and further in view of Dziadosz et al. (hereinafter Dziadosz), U.S. Patent 5,832,222 (see Applicant IDS).

52. In considering claim 33, Carter discloses a method of data storage management comprising the act of: providing a plurality of network-accessible storage devices capable of storing data, (col. 6, lines 7-12); implementing a plurality of storage management process instances, communicating storage messages between the storage management process instances, and identifying two or more storage devices at

Art Unit: 2151

different network locations associated with the data to be stored, (col. 3, lines 1-5, col. 4, lines 15-37, col. 6, lines 7-12, col. 7, lines 43-60); storing the data across two or more storage devices, (col. 23, lines 16-20); and, retrieving the stored data, (col. 7, lines 43-60).

Although the disclosed system of Carter shows substantial features of the claimed invention, it fails to expressly disclose: using parity information to verify the correctness of stored data.

Nevertheless, using parity to verify the correctness of stored data was well known in the art at the time of the present invention. In a similar field of endeavor, Litwin discloses a method and system for data recovery comprising: generating parity files for data to be stored on network-accessible devices, (col. 1, lines 17-20); verifying the correctness of the stored data using the parity data, (col. 3, lines 18-26), and retrieving a correct version of the data using the parity data, (col. 3, lines 51-57).

Thus, given the teachings of Litwin it would have been apparent to one of ordinary skill to modify the teachings of Carter to show using parity information to verify the correctness of stored data. This would show that there is a secure means for correcting errors and recovering data in the network accessible devices. This also would also further assure that the method disclosed by Carter is a fault tolerant method for preserving data transmitted to the network accessible devices, Litwin, col. 3, lines 51-57.

Although the modified teachings of Carter show substantial features of the claimed invention, they fail to expressly disclose: using a RAID-type distribution across two or more storage devices.

Nevertheless, the teachings of Carter provide, and suggest, a means for implementing a RAID-type distribution across the plurality of network devices, where Carter discloses: the storage management processes maintaining data coherence among the network nodes, and automatically replicating data for redundancy and fault tolerance, (col. 8, lines 42-50). Furthermore, RAID-type systems were well known in the art at the time of the present invention. This is demonstrated by Dziadosz in a similar field of endeavor wherein Dziadosz teaches: using a RAID-type distribution across two or more storage devices, (col. 3, lines 50-66).

Thus, if not implicit in the teachings of Carter, it would have been obvious to a person of ordinary skill in the art at the time of the present invention to further modify the teachings of Carter to show the storage management processes implementing a RAID-type distribution across the plurality of network accessible devices. Doing so would have improved performance of the storage management system taught by Carter, and would have also increased fault tolerance, (Dziadosz, col. 3, lines 50-66).

53. In considering claim 34, the disclosed system of Litwin teaches the parity record comprising data capable of correcting errors on another network-accessible storage device. See Litwin, col. 7, lines 56-60. The motivation to combine the



teachings of Carter and Litwin would be the same as that mentioned in consideration of claim 33.

54. In considering claim 35, the system disclosed by Carter provides a means for the parity data to comprise a mirror copy of the data to be stored. See col. 23, lines 16-20.

55. In considering claim 36, the system disclosed by Carter provides a means for the parity data to be stored in a single network storage node, and the data to be stored in two or more network storage nodes. See col. 23, lines 16-20.

56. In considering claim 37, the system disclosed by Carter provides a means for the parity data to be distributed across multiple storage nodes. See col. 8, lines 42-50.

57. In considering claim 39, the system disclosed by Carter further teaches: attempting to retrieve the stored data, detecting unavailability of one of the one or more network storage devices, and in response to detecting unavailability, reconstructing the correct data using a reconciliation log, (col. 36, lines 10-63).

Although the disclosed system of Carter shows substantial features of the claimed invention, it fails to expressly disclose: reconstructing the data using parity data.

Nevertheless, in a similar field of endeavor, Litwin discloses a method and system for data recovery comprising: reconstructing data in a failed data bucket using parity data, (col. 3, lines 18-26).

Given the teachings of Litwin it would have been apparent to one of ordinary skill to modify the teachings of Carter to show reconstructing data in a network storage device, after detecting the network storage device was unavailable, using parity data. This would show a secure means for correcting errors and recovering data in the network accessible devices. This also would also further assure that the method disclosed by Carter is a fault tolerant method for preserving data transmitted to the network accessible devices, Litwin, col. 3, lines 51-57.

58. In considering claim 40, the system disclosed by Carter provides a means for the act of storing data to comprise distributing non-identical but logically equivalent data in a storage node. See col. 8, lines 42-50.

59. In considering claim 41, the system disclosed by Carter provides a means for storing lossy equivalent data in a storage node. See col. 8, lines 42-50.

60. Claims 42-43, are rejected under 35 U.S.C. 103(a) as being unpatentable over Carter in view of Gershman et al. (hereinafter Gershman), U.S. Patent 6,199,099.

61. In considering claim 42, Carter discloses a method of data storage management comprising the acts of: providing a plurality of network-accessible storage devices capable of storing data, (col. 6, lines 37-43); implementing a plurality of storage management process instances to communicate with each other, and to store data to the network-accessible devices (col. 7, lines 43-60); changing the storage capacity under the control of the storage management processes without affecting accessibility of the data storage, (col. 9, lines 10-17).

Although the disclosed method of Carter shows substantial features of the claimed invention, it fails to explicitly disclose: the network devices transmitting heartbeat messages to be processed by the storage management processes.

Nevertheless, using heartbeat messages in network managing operations was well known in the art at the time of the present invention. In a similar field of endeavor, Gershman shows this. More specifically, Gershman teaches a system and method for managing a mobile communication network comprising: monitoring a network device by having the network device transmit heartbeat messages, (col. 47, lines 61-65).

Thus, given the teachings of Gershman, it would have been obvious to one of ordinary skill in the art to modify the teachings of Carter to show monitoring the data storage for faults by means of the plurality of storage management processes, wherein the monitoring comprises at least a portion of the plurality of network accessible storage devices transmitting heartbeat messages. Doing so would have provided an efficient means for the storage management processes to take appropriate action when a fault is detected through the monitoring, Gershman, col. 47, lines 65-67, col. 48, lines 1-4.

62. In considering claim 43, the method disclosed by Carter teaches compensating for faults by manipulating the data storage under control of the storage management processes. See col. 23, lines 12-27.

63. Claim 51 is rejected under 35 U.S.C. 103(a) as being unpatentable over Carter in view of McClain, U.S. patent 5,794,254 (supplied by applicant).

64. In considering claim 51, the disclosed method of Carter further teaches: security mechanisms when communicating, (col. 4, lines 38-50).

Although the disclosed system of Carter shows substantial features of the claimed invention, it fails to expressly disclose: encrypting storage messages before communicating.

Nevertheless, encrypting messages before communicating was well known in the art at the time of the present invention. This is exemplified in a similar field of endeavor where McClain discloses a method and system for backing up computer files at a remote site comprising: encrypting a storage message before communicating, (col. 6, lines 48-53).

Given the teachings of McClain it would have been apparent to one of ordinary skill to modify the teachings of Carter to show encrypting storage messages before communicating. This would provide a secure and safe means for storing data over a network, while preventing the data from being read by unauthorized individuals.

65. Claim 54, is rejected under 35 U.S.C. 103(a) as being unpatentable over Carter in view of Thompson, U.S. Patent 4,814,984.

66. In considering claim 54, although the disclosed method of Carter shows substantial features of the claimed invention, it fails to explicitly disclose: communicating an operational state between devices.

Nevertheless, communicating an operational state between devices was well known in the art at the time of the present invention. In a similar field of endeavor, Thompson shows this. More specifically, Thompson teaches a system and method for communicating between devices on a network comprising: maintaining a state table at each device on the network that indicates the operational state of other devices on the network, (col. 3, lines 29-33).

Thus, given the teachings of Thompson, it would have been obvious to one of ordinary skill in the art to modify the teachings of Carter to show communicating state information between the devices on the network, wherein the state information comprises access speed, transfer rate, network locality, physical locality, interconnectedness, security, reliability, political domain, capacity, or cost. This would have provided an efficient means for communicating valuable information between devices that would help in determining whether communication between the devices is appropriate, Thompson, col. 11, lines 26-30.

***Conclusion***

67. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire **THREE MONTHS** from the mailing date of this action. In the event a first reply is filed within **TWO MONTHS** of the mailing date of this final action and the advisory action is not mailed until after the end of the **THREE-MONTH** shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than **SIX MONTHS** from the date of this final action.

68. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Hassan Phillips whose telephone number is (571) 272-3940. The examiner can normally be reached on M-F 8:00am-5:00pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Zarni Maung can be reached on (571) 272-3939. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2151

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

HP/  
10/27/05

  
**ZARNI MAUNG**  
SUPERVISORY PATENT EXAMINER